

χ^2 and classical exact tests often wildly misreport significance; the remedy lies in computers

William Perkins*, Mark Tygert[†], and Rachel Ward[‡]

September 16, 2011

Abstract

If a discrete probability distribution in a model being tested for goodness-of-fit is not close to uniform, then forming the Pearson χ^2 statistic can involve division by nearly zero. This often leads to serious trouble in practice — even in the absence of round-off errors — as the present article illustrates via numerous examples. Fortunately, with the now widespread availability of computers, avoiding all the trouble is simple and easy: without the problematic division by nearly zero, the actual values taken by goodness-of-fit statistics are not humanly interpretable, but black-box computer programs can rapidly calculate their precise significance.

*Supported in part by NSF Grant OISE-0730136 and an NSF Postdoctoral Research Fellowship

[†]Supported in part by a Research Fellowship from the Alfred P. Sloan Foundation

[‡]Supported in part by an NSF Postdoctoral Research Fellowship and a Donald D. Harrington Faculty Fellowship

Contents

1	Introduction	3
2	Definitions of the test statistics	5
3	Hypothesis tests with parameter estimation	6
4	Data analysis	9
4.1	Synthetic examples	9
4.2	Zipf's power law of word frequencies	10
4.3	A Poisson law for radioactive decays	15
4.4	A Poisson law for counting with a hæmacytometer	15
4.5	A Hardy-Weinberg law for Rhesus blood groups	17
4.6	Symmetry between the self-reported health assessments of foreign- and US-born Asian Americans	20
4.7	A modified geometric law for the species of butterflies	23
4.8	A modified geometric law for religious affiliations	23
5	The power and efficiency of the root-mean-square	27
5.1	Examples without parameter estimation	28
5.1.1	A simple, illustrative example	28
5.1.2	Truncated power-laws	30
5.1.3	Additional truncated power-laws	31
5.1.4	Additional truncated power-laws, reversed	32
5.1.5	A final example with fully specified truncated power-laws	33
5.1.6	Modified Poisson distributions	34
5.1.7	A truncated power-law and a truncated geometric distribution	35
5.2	Examples with parameter estimation	36
5.2.1	A truncated power-law and a truncated geometric distribution	36
5.2.2	A rebinned geometric distribution and a truncated power-law	37
5.2.3	Truncated shifted Poisson distributions	38
5.2.4	Modified Poisson distributions	39
5.2.5	An example with a uniform tail	40
5.2.6	Another example with a uniform tail	41
5.2.7	A model with an integer-valued parameter	42
5.2.8	Truncated power-laws parameterized with a permutation	43
5.2.9	Another model parameterized with a permutation	44
5.2.10	A model with two parameters	45
A	Additional plots of power and efficiency	46
B	Convergence to asymptotic levels	56
	References	62

1 Introduction

A basic task in statistics is to ascertain whether a given set of independent and identically distributed (i.i.d.) draws does not come from a given “model,” where the model may consist of either a single fully specified probability distribution or a parameterized family of probability distributions. The present paper concerns the case in which the draws are discrete random variables, taking values in a finite or countable set. In accordance with the standard terminology, we will refer to the possible values of the discrete random variables as “bins” (“categories,” “cells,” and “classes” are common synonyms for “bins”).

A natural approach to ascertaining whether the i.i.d. draws do not come from the model uses a root-mean-square statistic. To construct this statistic, we estimate the probability distribution over the bins using the given i.i.d. draws, and then measure the root-mean-square difference between this empirical distribution and the model distribution; see, for example, [17], page 123 of [21], or Section 2 below. If the draws do in fact arise from the model, then with high probability this root-mean-square is not large. Thus, if the root-mean-square statistic is large, then we can be confident that the draws do not arise from the model.

To quantify “large” and “confident,” let us denote by x the value of the root-mean-square for the given i.i.d. draws; let us denote by X the root-mean-square statistic constructed for different i.i.d. draws that definitely do in fact come from the model (if the model is parameterized, then we draw from the distribution corresponding to the parameter given by a maximum-likelihood estimate for the experimental data). The significance level α is then defined to be the probability that $X \geq x$ (viewing X — but not x — as a random variable). The confidence level that the given i.i.d. draws do not arise from the model is the complement of the significance level, namely $1 - \alpha$. (See Remark 1.2 concerning our use of the term “significance level” as synonymous with the alternative term “ p -value.”)

Now, the significance levels for the simple root-mean-square statistic can be different functions of x for different model probability distributions. To avoid this seeming inconvenience asymptotically (in the limit of large numbers of draws), K. Pearson replaced the uniformly weighted mean in the root-mean-square with a weighted average; the weights are the reciprocals of the model probabilities associated with the various bins. This produces the classic χ^2 statistic — see, for example, [14] or formula (2) below. However, when model probabilities can be small (relative to others in the same distribution), this weighted average can involve division by nearly zero. As demonstrated below, dividing by nearly zero severely restricts the statistical power of χ^2 — even in the absence of round-off errors — especially when dividing by nearly zero for each of many bins. Moreover, this problem arises whether or not every bin contains several draws (see Remark 1.1).

The main thesis of the present article is that using only the classic χ^2 statistic is no longer appropriate, that certain alternatives are far superior now that computers are widely available. We demonstrate below that the simple root-mean-square, used in conjunction with the log-likelihood-ratio “ G^2 ” goodness-of-fit statistic, is generally preferable to the classic χ^2 statistic. (The log-likelihood-ratio also involves division by nearly zero, but tempers this somewhat by taking a logarithm.) We do not make any claim that this is the best possible alternative. In fact, the discrete Kolmogorov-Smirnov statistic (or one of its variants, such as the discrete Kuiper statistic — see, for example, [3] or [5]) can be more powerful than the root-mean-square in certain circumstances; in any case, the discrete Kolmogorov-

Smirnov statistic and the root-mean-square are similar in many ways, and complementary in others. We focus on the root-mean-square largely because it is so simple and easy to understand; for example, computing the confidence levels of the root-mean-square in the limit of large numbers of draws is trivial, even when estimating continuous parameters via maximum-likelihood methods (see [15] and [16]). Furthermore, the classic χ^2 statistic is just a weighted version of the root-mean-square, facilitating their comparison. Finally, χ^2 and the root-mean-square coincide when the model distribution is uniform.

Please note that all statistical tests reported in the present paper (including those involving the χ^2 statistic) are exact; we compute significance levels via Monte-Carlo simulations providing guaranteed error bounds (see Section 3 below). In all numerical results reported below, we generated random numbers via the C programming language procedure given on page 9 of [13], implementing the recommended complementary multiply with carry.

To be sure, the problem with χ^2 is neither subtle nor esoteric. For a particularly revealing example, see Subsection 4.5 below.

Appropriate rebinning to uniformize the probabilities associated with the bins can mitigate much of the problem with χ^2 . Yet rebinning is a black art that is liable to improperly influence the result of a goodness-of-fit test. Moreover, rebinning requires careful extra work, making χ^2 less easy-to-use. A principal advantage of the root-mean-square is that it does not require any rebinning; indeed, the root-mean-square is most powerful without any rebinning.

Remark 1.1. In many of our examples, there is a bin for which the expected number of draws is very small under the model. Please note that, although it is natural for the expected numbers of draws for some bins to be very small, especially when the model has many bins, the advantage of the root-mean-square over χ^2 is substantial even when the expected number of draws is at least five for every bin; see, for example, Subsection 5.1.1 or Subsection 5.2.6.

Remark 1.2. Please beware that we treat “significance level” as synonymous with the alternative term “ p -value.” These two terms are not exactly the same in the classical terminology. However, the older concept of “significance level” is no longer very relevant, due to the proliferation of computer technology; there is no longer much reason to calculate and store tables of thresholds for goodness-of-fit statistics at arbitrarily fixed significance levels — we can now compute “ p -values” on the fly, as needed. The objective of a significance test is not really to accept or reject a hypothesis at some arbitrary threshold of significance, but instead to provide significance levels that can inform statisticians’ further analysis.

Remark 1.3. Goodness-of-fit tests are probably most useful in practice not for ascertaining whether a model is correct or not, but for determining whether the discrepancy between the model and experiment is larger than expected random fluctuations. While models outside the physical sciences typically are not exactly correct, testing the validity of using a model for virtually any purpose requires knowing whether observed discrepancies are due to inaccuracies or inadequacies in the models or (on the contrary) could be due to chance arising from necessarily finite sample sizes. Thus, goodness-of-fit tests are critical even when the models are not supposed to be exactly correct, in order to gauge the size of the unavoidable random fluctuations. For further clarification, see [10] and the remarkably extensive title of the original article [14] that introduced the χ^2 test for goodness-of-fit.

Remark 1.4. Combining the root-mean-square methodology and the statistical bootstrap (see, for example, [6]) should produce a test for whether two separate sets of draws arise from the same or from different distributions, when each set is taken i.i.d. from some (unspecified) distribution; the two distributions associated with the sets may differ. This is related to testing for association/independence in contingency-tables/cross-tabulations that have only two rows.

2 Definitions of the test statistics

In this section, we review the definitions of four goodness-of-fit statistics — the root-mean-square, χ^2 , the log-likelihood-ratio or G^2 , and the Freeman-Tukey or Hellinger distance. The latter three statistics are the best-known members of the standard Cressie-Read power-divergence family (see, for example, [17]). We use $p_1, p_2, \dots, p_{n-1}, p_n$ to denote the expected fractions of m i.i.d. draws falling in n bins, numbered $1, 2, \dots, n-1, n$, respectively, and we use $q_1, q_2, \dots, q_{n-1}, q_n$ to denote the observed fractions of the m draws falling in the respective bins. That is, $p_1, p_2, \dots, p_{n-1}, p_n$ are the probabilities associated with the respective bins in the *model* distribution, whereas $q_1, q_2, \dots, q_{n-1}, q_n$ are the fractions of the m draws falling in the respective bins when we take the draws from a distribution that may differ from the model — their *actual* distribution. Specifically, if $i_1, i_2, \dots, i_{m-1}, i_m$ are the observed i.i.d. draws, then q_k is $\frac{1}{m}$ times the number of $i_1, i_2, \dots, i_{m-1}, i_m$ falling in bin k , for $k = 1, 2, \dots, n-1, n$. If the model is parameterized by a parameter θ , then the probabilities $p_1, p_2, \dots, p_{n-1}, p_n$ are functions of θ ; if the model is fully specified, then we can view the probabilities $p_1, p_2, \dots, p_{n-1}, p_n$ as constant as functions of θ . We use $\hat{\theta}$ to denote a maximum-likelihood estimate of θ obtained from $q_1, q_2, \dots, q_{n-1}, q_n$.

With this notation, the root-mean-square statistic is

$$X = \sqrt{\frac{1}{n} \sum_{k=1}^n (q_k - p_k(\hat{\theta}))^2}. \quad (1)$$

We use the designation “root-mean-square” to refer to X .

The classical Pearson χ^2 statistic is

$$\chi^2 = m \sum_{k=1}^n \frac{(q_k - p_k(\hat{\theta}))^2}{p_k(\hat{\theta})}, \quad (2)$$

under the convention that $(q_k - p_k(\hat{\theta}))^2/p_k(\hat{\theta}) = 0$ if $p_k(\hat{\theta}) = 0 = q_k$. We use the standard designation “ χ^2 ” to refer to χ^2 .

The log-likelihood-ratio or “ G^2 ” statistic is

$$G^2 = 2m \sum_{k=1}^n q_k \ln \left(\frac{q_k}{p_k(\hat{\theta})} \right), \quad (3)$$

under the convention that $q_k \ln(q_k/p_k(\hat{\theta})) = 0$ if $q_k = 0$. We use the common designation “ G^2 ” to refer to G^2 .

The Freeman-Tukey or Hellinger-distance statistic is

$$H^2 = 4m \sum_{k=1}^n \left(\sqrt{q_k} - \sqrt{p_k(\hat{\theta})} \right)^2 = 4m \sum_{k=1}^n \left[(q_k - p_k(\hat{\theta}))^2 / \left(\sqrt{q_k} + \sqrt{p_k(\hat{\theta})} \right)^2 \right]. \quad (4)$$

We use the well-known designation “Freeman-Tukey” to refer to H^2 .

In the limit that the number m of draws is large, the distributions of χ^2 defined in (2), G^2 defined in (3), and H^2 defined in (4) are all the same when the actual underlying distribution of the draws comes from the model (see, for example, [17]). However, when the number m of draws is not large, then their distributions can differ substantially. In all our data and power analyses, we compute confidence levels via Monte-Carlo simulations, without relying on the number m of draws to be large.

3 Hypothesis tests with parameter estimation

In this section, we discuss the testing of hypotheses involving parameterized models: Given a family $p(\theta)$ of probability distributions parameterized by θ , and given observed i.i.d. draws from some actual underlying (unknown) distribution \tilde{p} , we would like to test the hypothesis

$$H'_0 : \text{for some } \theta, \tilde{p} = p(\theta), \quad (5)$$

against the alternative

$$H'_1 : \text{for all } \theta, \tilde{p} \neq p(\theta). \quad (6)$$

Given only finitely many draws, the significance level for such a test would have to be independent of the parameter θ , since the proper value for θ is unknown (θ is known as a nuisance parameter). Unfortunately, it is not clear how to devise such a test when the probability distributions are discrete. None of the standard methods (including χ^2 , the log-likelihood-ratio, the Freeman-Tukey/Hellinger distance, and other power-divergence statistics) produce significance levels that are independent of the parameter θ . Some methods do produce significance levels that are independent of θ in the limit of large numbers of draws, but this is not especially useful, since in the limit of large numbers of draws any actual parameter θ would be almost surely known anyway (see Appendix B for further elaboration).

In the present paper, we test the significance of assuming

$$H_0 : \tilde{p} = p(\hat{\theta}) \text{ for the particular observed value of } \hat{\theta}, \quad (7)$$

where $\hat{\theta}$ is a maximum-likelihood estimate of θ ; that is, H_0 is the hypothesis that $\tilde{p} = p(\hat{\theta})$ for the value of $\hat{\theta}$ associated with the single realization of the experiment that was measured (subsequent repetitions of the experiment, including those considered when calculating the significance level as in Remark 3.3, can yield different estimates of the parameter, even though the repetitions’ actual distribution \tilde{p} is the same). Of course, the accuracy of the estimate $\hat{\theta}$ generally improves as the number of draws increases; in fact, testing (5) and testing (7) are asymptotically equivalent, in the limit of large numbers of draws (see [16]).

As testing the hypothesis H'_0 defined in (5) does not seem to be feasible in general when the probability distributions are discrete and there are more than just a few bins, we focus on testing the closely related assumption H_0 defined in (7). The latter is more relevant

for many applications, anyways — plots typically display the particular fitted distribution in (7); interpreting such plots naturally involves (7). All tests of the present paper concern the significance of assuming H_0 defined in (7) (if the model is fully specified, then the probability distribution $p(\theta)$ is the same for all θ). Please be sure to bear in mind Remark 1.3 of Section 1.

Remark 3.1. Another means of handling nuisance parameters is to test the hypothesis

$$H_0'' : \tilde{p} = p(\hat{\theta}) \text{ for all possible realizations of the experiment;} \quad (8)$$

that is, H_0'' is the hypothesis that $\tilde{p} = p(\hat{\theta})$ and that $p(\hat{\theta})$ always takes exactly the same value during repetitions of the experiment. The assumption that (8) is true seems to be more extreme, a more substantial departure from (5), than (7). Nevertheless, testing (8) is standard; see, for example, Section 6 of [4]. Assuming (8) amounts to conditioning (5) on a statistic that is minimally sufficient for estimating θ ; computing the associated significance levels is not always trivial. Testing the significance of assuming (7) would seem to be more apropos in practice for applications in which the experimental design does not enforce that repeated experiments always yield the same value for $p(\hat{\theta})$.

Remark 3.2. The parameter θ can be integer-valued, real-valued, complex-valued, vector-valued, matrix-valued, or any combination of the many possibilities. For instance, when we do not know the proper ordering of the bins a priori, we must include a parameter that contains a permutation (or permutation matrix) specifying the order of the bins; maximum-likelihood estimation then entails sorting the model and all empirical frequencies (whether experimental or simulated) — see Subsection 4.2 for details. With Remark 3.3, we need not contemplate how many degrees of freedom are in a permutation.

Remark 3.3. To compute the level of significance of assuming (7), we can use Monte-Carlo simulations (very similar to those in [3]). First, we estimate the parameter θ from the m given experimental draws, obtaining $\hat{\theta}$, and then calculate the statistic under consideration (χ^2 , G^2 , Freeman-Tukey, or the root-mean-square), using the given data and taking the model distribution to be $p(\hat{\theta})$. We then run many simulations. To conduct a single simulation, we perform the following three-step procedure:

1. we generate m i.i.d. draws according to the model distribution $p(\hat{\theta})$, where $\hat{\theta}$ is the estimate calculated from the experimental data,
2. we estimate the parameter θ from the data generated in Step 1, obtaining a new estimate $\tilde{\theta}$, and
3. we calculate the statistic under consideration (χ^2 , G^2 , Freeman-Tukey, or the root-mean-square), using the data generated in Step 1 and taking the model distribution to be $p(\tilde{\theta})$, where $\tilde{\theta}$ is the estimate calculated in Step 2 from the data generated in Step 1.

After conducting many such simulations, we may estimate the confidence level for rejecting (7) as the fraction of the statistics calculated in Step 3 that are less than the statistic calculated from the empirical data. (Recall that a significance level of α is the same as a confidence level of $1 - \alpha$.) The accuracy of the estimated confidence level is inversely proportional to the square root of the number of simulations conducted; for details, see Remark 3.4 below. This procedure works since, by definition, the confidence level is the probability that

$$d \left[\begin{pmatrix} Q_1 \\ Q_2 \\ \vdots \\ Q_{n-1} \\ Q_n \end{pmatrix}, \begin{pmatrix} p_1(\Theta) \\ p_2(\Theta) \\ \vdots \\ p_{n-1}(\Theta) \\ p_n(\Theta) \end{pmatrix} \right] < d \left[\begin{pmatrix} q_1 \\ q_2 \\ \vdots \\ q_{n-1} \\ q_n \end{pmatrix}, \begin{pmatrix} p_1(\hat{\theta}) \\ p_2(\hat{\theta}) \\ \vdots \\ p_{n-1}(\hat{\theta}) \\ p_n(\hat{\theta}) \end{pmatrix} \right], \quad (9)$$

where

- n is the number of all possible values that the draws can take,
- d is the measure of the discrepancy between two probability distributions over n bins (i.e., between two vectors each with n entries) that is associated with the statistic under consideration (d is the Euclidean distance for the root-mean-square, a weighted Euclidean distance for χ^2 , the Hellinger distance for the Freeman-Tukey statistic, and the relative entropy — the Kullback-Leibler divergence — for the log-likelihood-ratio),
- $q_1, q_2, \dots, q_{n-1}, q_n$ are the fractions of the m given experimental draws falling in the respective bins,
- $\hat{\theta}$ is the estimate of θ obtained from $q_1, q_2, \dots, q_{n-1}, q_n$,
- $Q_1, Q_2, \dots, Q_{n-1}, Q_n$ are the fractions of m i.i.d. draws falling in the respective bins when taking the draws from the distribution $p(\hat{\theta})$ assumed in (7), and
- Θ is the estimate of the parameter θ obtained from $Q_1, Q_2, \dots, Q_{n-1}, Q_n$ (note that Θ is not necessarily always equal to $\hat{\theta}$: even under the null hypothesis, repetitions of the experiment could yield different estimates of the parameter; see also Remark B.2).

When taking the probability that (9) occurs, only the left-hand side is random — we regard the left-hand side of (9) as a random variable and the right-hand side as a fixed number determined via the experimental data. As with any probability, to compute the probability that (9) occurs, we can calculate many independent realizations of the random variable and observe that the fraction which satisfy (9) is a good approximation to the probability when the number of realizations is large; Remark 3.4 details the accuracy of the approximation. (The procedure in the present remark follows this prescription to estimate confidence levels.)

Remark 3.4. The standard error of the estimate from Remark 3.3 for an exact significance level of α is $\sqrt{\alpha(1-\alpha)/\ell}$, where ℓ is the number of Monte-Carlo simulations conducted to produce the estimate. Indeed, each simulation has probability α of producing a statistic that is greater than or equal to the statistic corresponding to an exact significance level of α . Since the simulations are all independent, the number of the ℓ simulations that produce statistics greater than or equal to that corresponding to level α follows the binomial distribution with ℓ trials and probability α of success in each trial. The standard deviation of the number of simulations whose statistics are greater than or equal to that corresponding to level α is therefore $\sqrt{\ell\alpha(1-\alpha)}$, and so the standard deviation of the *fraction* of the simulations producing such statistics is $\sqrt{\alpha(1-\alpha)/\ell}$. Of course, the fraction itself is the Monte-Carlo estimate of the exact significance level (we use this estimate in place of the unknown α when calculating the standard error $\sqrt{\alpha(1-\alpha)/\ell}$).

4 Data analysis

In this section, we use several data sets to investigate the performance of goodness-of-fit statistics. The root-mean-square generally performs much better than the classical statistics. We take the position that a user of statistics should not have to worry about rebinning; we discuss rebinning only briefly. We compute all significance levels via Monte Carlo as in Remark 3.3; Remark 3.4 details the guaranteed accuracy of the computed significance levels.

4.1 Synthetic examples

To better explicate the performance of the goodness-of-fit statistics, we first analyze some toy examples. We consider the model distribution

$$p_1 = \frac{1}{4}, \tag{10}$$

$$p_2 = \frac{1}{4}, \tag{11}$$

and

$$p_k = \frac{1}{2n-4} \tag{12}$$

for $k = 3, 4, \dots, n-1, n$. For the empirical distribution, we first use $m = 20$ draws, with 15 in the first bin, 5 in the second bin, and no draw in any other bin. This data is clearly unlikely to arise from the model specified in (10)–(12), but we would like to see exactly how well the various goodness-of-fit statistics detect the obvious discrepancy.

Figure 1 plots the significance levels for testing whether the empirical data arises from the model specified in (10)–(12). We computed the significance levels via 4,000,000 Monte-Carlo simulations (that is, 4,000,000 per empirical significance level being evaluated), with each simulation taking $m = 20$ draws from the model. The root-mean-square consistently and with extremely high confidence rejects the hypothesis that the data arises from the model, whereas the classical statistics find less and less evidence for rejecting the hypothesis as the number n of bins increases; in fact, the significance levels for the classical statistics get very close to 1 as n increases — the discrepancy of (12) from 0 is usually less than the discrepancy of (12) from a typical realization drawn from the model, since under the model the sum of the expected numbers of draws in bins $3, 4, \dots, n-1, n$ is $m/2$.

Figure 1 demonstrates that the root-mean-square can be much more powerful than the classical statistics, rejecting with nearly 100% confidence while the classical statistics report nearly 0% confidence for rejection. Moreover, the classical statistics can report significance levels very close to 1 even when the data manifestly does not arise from the model. (Incidentally, the model for smaller n can be viewed as a rebinning of the model for larger n . The classical statistics do reject the model for smaller n , while asserting for larger n that there is no evidence for rejecting the model.) The performance of the classical statistics displays a dramatic dependence on the number $(n-2)$ of unlikely bins in the model, even though the data are the same for all n . This suggests a sure-fire scheme for supporting any model (no matter how invalid) with arbitrarily high significance: just append enough irrelevant, more or less uniformly improbable bins to the model, and then report the significance levels




Figure 1: Significance levels for the hypothesis that the model (10)–(12) agrees with the data of 15 draws in the first bin, 5 draws in the second bin, and no draw in any other bin

for the classical goodness-of-fit statistics. In contrast, the root-mean-square robustly and reliably rejects the invalid model, independently of the size of the model.

We will see in the following section that the classic Zipf power law behaves similarly.

For another example, we again consider the model specified in (10)–(12). For the empirical distribution, we now use $m = 96$ draws, with 36 in the first bin, 12 in the second bin, 1 each for bins 3, 4, ..., 49, 50, and no draw in any other bin. As before, this data clearly is unlikely to arise from the model specified in (10)–(12), but we would like to see exactly how well the various goodness-of-fit statistics detect the obvious discrepancy.

Figure 2 plots the significance levels for testing whether the empirical data arises from the model specified in (10)–(12). We computed the significance levels via 160,000 Monte-Carlo simulations (that is, 160,000 per empirical significance level being evaluated), with each simulation taking $m = 96$ draws from the model. Yet again, the root-mean-square consistently and confidently rejects the hypothesis that the data arises from the model, whereas the classical statistics find little evidence for rejecting the manifestly invalid model.

4.2 Zipf’s power law of word frequencies

Zipf popularized his eponymous law by analyzing four “chief sources of statistical data referred to in the main text [23]” (this is a quotation from the “Notes and References” section — page 311 — of [23]); in [23], the chief source for the English language is [7]. We revisit the data from [7] in the present subsection to assess the performance of the goodness-of-fit statistics.

We first analyze List 1 of [7], which consists of 2,890 different English words, such that




Figure 2: Significance levels for the hypothesis that the model (10)–(12) agrees with the data of 36 draws in the first bin, 12 draws in the second bin, 1 draw each in bins 3, 4, \dots , 49, 50, and no draw in any other bin

there are 13,825 words in total counting repetitions; the words come from the Buffalo Sunday News of August 8, 1909. We randomly choose $m = 10,000$ of the 13,825 words to obtain a corpus of $m = 10,000$ draws over 2,890 bins. Figure 3 plots the frequencies of the different words when sorted in rank order (so that the frequencies are nonincreasing). Using goodness-of-fit statistics we test the significance of the (null) hypothesis that the empirical draws actually arise from the Zipf distribution

$$p_k(\theta) = \frac{C_1}{\theta(k)} \quad (13)$$

for $k = 1, 2, \dots, n - 1, n$, where θ is a permutation of the integers $1, 2, \dots, n - 1, n$, and

$$C_1 = \frac{1}{\sum_{k=1}^n 1/k}; \quad (14)$$

we estimate the permutation θ via maximum-likelihood methods, that is, by sorting the frequencies: first we choose k_1 to be the number of a bin containing the greatest number of draws among all n bins, then we choose k_2 to be the number of a bin containing the greatest number of draws among the remaining $n - 1$ bins, then we choose k_3 to be the number of a bin containing the greatest among the remaining $n - 2$ bins, and so on, and finally we find θ such that $\theta(k_1) = 1, \theta(k_2) = 2, \dots, \theta(k_{n-1}) = n - 1, \theta(k_n) = n$. We have to obtain the ordering θ from the data via such sorting since we do not know the proper ordering a priori.

Similarly, we do not know the proper value of the number n of bins, so in Figure 4 we plot significance levels (each computed via 40,000 Monte-Carlo simulations) for varying values of

n ; although List 1 of [7] involves only 2,890 distinct words, we must also include bins for words that did not appear in the original list, words whose frequencies are zeros for List 1 of [7]. Note that Figure 4 displays the significance levels with $n = 2,890$ for reference, even though n must be independent of the data, and so n must be substantially larger than 2,890 in order for the assumptions of goodness-of-fit testing to hold.

With respect to testing goodness-of-fit, the number n of bins is the number of words in the dictionary from which List 1 of [7] was drawn. It is not clear a priori which dictionary is appropriate. Fortunately, the significance levels for the root-mean-square are always 0 to several digits of accuracy, independent of the value of n — the root-mean-square determines that List 1 does not follow the classic Zipf distribution (defined in (13) and (14)) for any n . In contrast, the significance levels for the classical statistics vary wildly depending on the value of n . In fact, for any of the classical statistics, and for any prescribed number α between 0.05 and 0.95, there is at least one value of n between 4,000 and 40,000 such that the significance level is α . Thus, without knowing the proper size of the dictionary a priori, the classical statistics are meaningless.

Unsurprisingly, analyzing List 5 of [7] produces results analogous to those reported above for List 1. List 5 consists of 6,002 different English words, such that there are 43,989 words in total counting repetitions; the words come from amalgamating Lists 1–4 of [7]. We randomly choose $m = 20,000$ of the 43,989 words to obtain a corpus of $m = 20,000$ draws over 6,002 bins. Figure 5 plots the frequencies of the different words when sorted in rank order (so that the frequencies are nonincreasing).

Again we do not know the proper value of the number n of bins, so in Figure 6 we plot significance levels (each computed via 40,000 Monte-Carlo simulations) for varying values of n ; although List 5 of [7] involves only 6,002 distinct words, we must also include bins for words that did not appear in the original list, words whose frequencies are zeros for List 5 of [7]. Please note that Figure 6 displays the significance levels with $n = 6,002$ for reference, even though n must be independent of the data, and so n must be substantially larger than 6,002 in order for the assumptions of goodness-of-fit testing to hold. Comparing Figures 4 and 6 shows that the above remarks about List 1 pertain to the analysis of the larger List 5, too. Once again, without knowing the proper size of the dictionary a priori, the classical statistics are meaningless, whereas the root-mean-square is very powerful.

Interestingly, by introducing parameters θ_1 , θ_2 , and θ_3 to fit perfectly the bins containing the three greatest numbers of draws, a truncated power-law becomes a good fit for the corpus of 20,000 words drawn randomly from List 5 of [7], with the number n of bins set to 7,500. Indeed, let us consider the model

$$p_k(\theta_0, \theta_1, \theta_2, \theta_3, \theta_4) = \begin{cases} \theta_1, & \theta_0(k) = 1 \\ \theta_2, & \theta_0(k) = 2 \\ \theta_3, & \theta_0(k) = 3 \\ C/(\theta_0(k))^{\theta_4}, & \theta_0(k) = 4, 5, \dots, 7499, 7500 \end{cases}, \quad (15)$$

where

$$C = C_{\theta_1, \theta_2, \theta_3, \theta_4} = \frac{1 - \theta_1 - \theta_2 - \theta_3}{\sum_{k=4}^{7500} 1/k^{\theta_4}}, \quad (16)$$

with θ_0 being a permutation of the integers 1, 2, \dots , 7499, 7500, and θ_1 , θ_2 , θ_3 , θ_4 being nonnegative real numbers; we estimate θ_0 , θ_1 , θ_2 , θ_3 , θ_4 via maximum-likelihood methods,




Figure 3: Numbers of occurrences of the various words (one bin for each distinct word) in a corpus of 10,000 random draws from List 1 of [7]

determining θ_0 by sorting as discussed above, and setting θ_1 , θ_2 , and θ_3 to be the three greatest relative frequencies. This model fits the empirical data exactly in the bins whose probabilities under the model are θ_1 , θ_2 , and θ_3 — there will be no discrepancy between the data and the model in those bins — so that these bins do not contribute to any goodness-of-fit statistic, aside from altering the number of draws in the remaining bins. Of the 20,000 total draws in the given experimental data, 16,486 do not fall in the bins associated with the three most frequently occurring words. The maximum-likelihood estimate of the power-law exponent θ_4 for the experimental data turns out to be about 1.0484.

For the model defined in (15) and (16), the significance levels calculated via 4,000,000 Monte-Carlo simulations are

- χ^2 : .510
- G^2 : .998
- Freeman-Tukey: 1.000
- root-mean-square: .587

Thus, all four statistics indicate that the truncated power-law model defined in (15) and (16) is a good fit. This is in accord with Figure 5, in which all but the three greatest frequencies appear to follow a truncated power-law.




Figure 4: Significance levels for the data plotted in Figure 3 to follow the Zipf distribution




Figure 5: Numbers of occurrences of the various words (one bin for each distinct word) in a corpus of 20,000 random draws from List 5 of [7]




Figure 6: Significance levels for the data plotted in Figure 5 to follow the Zipf distribution

4.3 A Poisson law for radioactive decays

Table 1 summarizes the classic example of a Poisson-distributed experiment in radioactive decay from [19]; Figure 7 plots the data, along with the Poisson distribution whose mean is the same as the data's. Figure 8 reports the significance levels for testing whether the data, while retaining only bins $1, 2, \dots, n-1, n$, are distributed according to a Poisson distribution (the model Poisson distribution is also truncated to the first n bins, with the mean estimated from the data). Since the total number m of draws depends little on the numbers in bins $13, 14, 15, \dots$, the truncation amounts to ignoring draws in bins $n+1, n+2, n+3, \dots$ when $n \geq 12$, and demonstrates that the scant experimental draws in bins 13–15 strongly influence the significance levels of the classical statistics. We computed the significance levels via 40,000 Monte-Carlo simulations (for each number n of bins and each of the four statistics), estimating the mean of the model Poisson distribution for each simulated data set. All four goodness-of-fit statistics indicate reasonably good agreement between the data and a Poisson distribution; the classical statistics are very sensitive in the tail to discrepancies between the data and the model distribution, whereas the root-mean-square is relatively insensitive to the truncation after 12 or more bins.

4.4 A Poisson law for counting with a hæmacytometer

Page 357 of [20] reports the number of yeast cells observed in each of 400 squares in a hæmacytometer microscope slide. Table 2 displays the counts; Figure 9 plots them, along with the Poisson distribution whose mean matches the data's. The significance levels for the data to arise from a Poisson distribution (with the mean estimated from the data) are

Table 1: Numbers of α -particles emitted by a film of polonium in 2608 intervals of 7.5 seconds

bin number	number of particles observed	
	in an interval of 7.5 seconds	number of such intervals
1	0	57
2	1	203
3	2	383
4	3	525
5	4	532
6	5	408
7	6	273
8	7	139
9	8	45
10	9	27
11	10	10
12	11	4
13	12	0
14	13	1
15	14	1
16, 17, 18, ...	15, 16, 17, ...	0
1, 2, 3, 4, 5, ...	0, 1, 2, 3, 4, ...	2608




Figure 7: The data in Table 1 (the dots) and the best-fit Poisson distribution (the lines)




Figure 8: Significance levels for the distribution of Table 1 to be Poisson

- χ^2 : .627
- G^2 : .365
- Freeman-Tukey: .111
- root-mean-square: .490

We calculated the significance levels via 4,000,000 Monte-Carlo simulations, estimating the mean of the model Poisson distribution for each simulated data set. Evidently, all four statistics report that a Poisson distribution is a reasonably good model for the experimental data.

4.5 A Hardy-Weinberg law for Rhesus blood groups

In a population with suitably random mating, the proportions of pairs of Rhesus haplotypes in members of the population (each member has one pair) can be expected to follow the Hardy-Weinberg law (see, for example, [11]), namely to arise via random sampling from the model

$$p_{j,k}(\theta_1, \theta_2, \dots, \theta_8, \theta_9) = \begin{cases} 2 \cdot \theta_j \cdot \theta_k, & j > k \\ (\theta_k)^2, & j = k \end{cases} \quad (17)$$

for $j, k = 1, 2, \dots, 8, 9$ with $j \geq k$, under the constraint that

$$\sum_{k=1}^9 \theta_k = 1, \quad (18)$$

Table 2: Numbers of yeast cells in 400 squares of a hæmacytometer

bin number	number of yeast in a square	number of such squares
1	0	0
2	1	20
3	2	43
4	3	53
5	4	86
6	5	70
7	6	54
8	7	37
9	8	18
10	9	10
11	10	5
12	11	2
13	12	2
14, 15, 16, ...	13, 14, 15, ...	0
1, 2, 3, 4, 5, ...	0, 1, 2, 3, 4, ...	400




Figure 9: The data in Table 2 (the dots) and the best-fit Poisson distribution (the lines)

Table 3: Frequencies of pairs of Rhesus haplotypes

		k								
$j \backslash k$		1	2	3	4	5	6	7	8	9
1		1236								
2		120	3							
3		18	0	0						
4		982	55	7	249					
5		32	1	0	12	0				
6		2582	132	20	1162	29	1312			
7		6	0	0	4	0	4	0		
8		2	0	0	0	0	0	0	0	
9		115	5	2	53	1	149	0	0	4

where the parameters $\theta_1, \theta_2, \dots, \theta_8, \theta_9$ are the proportions of the nine Rhesus haplotypes in the population (their maximum-likelihood estimates are the proportions of the haplotypes in the given data). For $j, k = 1, 2, \dots, 8, 9$ with $j \geq k$, therefore, $p_{j,k}$ is the expected probability that the pair of haplotypes in the genome of an individual is the pair j and k .

In this formulation, the hypothesis of suitably random mating entails that the members of the sample population are i.i.d. draws from the model specified in (17); if a goodness-of-fit statistic rejects the model with high confidence, then we can be confident that mating has not been suitably random. Table 3 provides data on $m = 8297$ individuals; we duplicated Figure 3 of [11] to obtain Table 3.

The significance levels calculated via 4,000,000 Monte-Carlo simulations are

- χ^2 : .693
- G^2 : .600
- Freeman-Tukey: .562
- negative log-likelihood (see Remark 4.2 below): .649
- root-mean-square: .039

Unlike the root-mean-square, the classical statistics are blind to the significant discrepancy between the data and the Hardy-Weinberg model.

Remark 4.1. For the example of the present subsection, rejecting the null hypothesis (5) from Section 3 might seem in principle to be more interesting than rejecting the assumption (7). Fortunately, the difference between (5) and (7) is essentially irrelevant for the root-mean-square in this example. Indeed, the root-mean-square is not very sensitive to bins associated with the parameters whose estimated values are potentially inaccurate — the potentially inaccurate estimates are all small, and the root-mean-square is not very sensitive to bins whose probabilities under the model are small relative to others.

Table 4: Frequencies of antigen genotypes

		k			
$j \backslash k$	1	2	3	4	
1	0				
2	3	1			
3	5	18	1		
4	3	7	5	2	

Remark 4.2. The term “negative log-likelihood” used in the present section refers to the statistic that is simply the negative of the logarithm of the likelihood. The negative log-likelihood is the same statistic used in the generalization of Fisher’s exact test discussed in [11]; unlike G^2 , this statistic involves only one likelihood, not the ratio of two. We mention the negative log-likelihood just to facilitate comparisons with [11]; we are not asserting that the likelihood on its own (rather than in a ratio) is a good gauge of the relative sizes of deviations from a model.

Remark 4.3. Table 4 provides data on $m = 45$ individuals from the other set of real-world measurements given in [11]; we duplicated Figure 2 of [11] to obtain Table 4. The associated Hardy-Weinberg model is then the same as (17), but with only four parameters, $\theta_1, \theta_2, \theta_3, \theta_4$, such that $\sum_{k=1}^4 \theta_k = 1$. The significance levels calculated via 4,000,000 Monte-Carlo simulations are

- χ^2 : .021
- G^2 : .013
- Freeman-Tukey: .027
- negative log-likelihood (see Remark 4.2 above): .016
- root-mean-square: .0019

Again the root-mean-square is more powerful than the classical statistics (though in this case all these statistics report significant discrepancies between the data and the Hardy-Weinberg model).

4.6 Symmetry between the self-reported health assessments of foreign- and US-born Asian Americans

Using propensity scores, [8] matched each of 335 surveyed foreign-born Asian Americans to a similar surveyed US-born Asian American. Table 5 duplicates Table 4 of [8], which tabulates the numbers of matched pairs reporting various combinations of self-rated physical health; the model used for generating the propensity scores did not explicitly incorporate the

health ratings. Table 5 does not reveal any significant difference between foreign-born Asian Americans' ratings of their health and US-born Asian Americans'. Indeed, the significance levels calculated via 4,000,000 Monte-Carlo simulations for testing the symmetry of Table 5 are

- χ^2 : .784
- G^2 : .739
- Freeman-Tukey: .642
- root-mean-square: .973

After noting that χ^2 does not reveal any statistically significant asymmetry in Table 5, [8] reports that, “to address the issue of power of this test, we investigated what is the smallest departure from symmetry that our test could detect...” Such an investigation requires considering modifications to Table 5. Table 6 provides one possible modification. The significance levels calculated via 4,000,000 Monte-Carlo simulations for testing the symmetry of Table 6 are

- χ^2 : .109
- G^2 : .123
- Freeman-Tukey: .155
- root-mean-square: .014

Evidently, the root-mean-square is more powerful for detecting the asymmetry of Table 6.

Table 7 provides another hypothetical cross-tabulation. The significance levels calculated via 64,000,000 Monte-Carlo simulations for testing the symmetry of Table 7 are

- χ^2 : .0015
- G^2 : .00016
- Freeman-Tukey: .000006, i.e., 6E-6
- root-mean-square: .131

The classical statistics are much more powerful for detecting the asymmetry of Table 7, contrasting how the root-mean-square is more powerful for detecting the asymmetry of Table 6. Indeed, the root-mean-square statistic is not very sensitive to relative discrepancies between the model and actual distributions in bins whose associated model probabilities are small. When sensitivity in these bins is desirable, we recommend using both the root-mean-square statistic and an asymptotically equivalent variation of χ^2 , such as the log-likelihood-ratio G^2 ; see, for example, [17].

Table 5: Self-reported physical health for matched pairs of Asian Americans

		foreign-born				
		excellent	very good	good	fair	poor
US-born	excellent	10	21	22	5	0
	very good	24	53	43	15	3
	good	21	43	34	11	0
	fair	3	11	8	4	1
	poor	1	1	1	0	0

Table 6: A variation on Table 5

		foreign-born				
		excellent	very good	good	fair	poor
US-born	excellent	10	21	22	5	0
	very good	24	53	<i>56</i>	15	3
	good	21	<i>30</i>	34	11	0
	fair	3	11	8	4	1
	poor	1	1	1	0	0

Table 7: Another variation on Table 5

		foreign-born				
		excellent	very good	good	fair	poor
US-born	excellent	10	21	22	5	0
	very good	24	53	43	15	3
	good	21	43	34	<i>19</i>	0
	fair	3	11	<i>0</i>	4	1
	poor	1	1	1	0	0

4.7 A modified geometric law for the species of butterflies

C. B. Williams, R. A. Fisher, and A. S. Corbet reported in [9] on 5300 butterflies from 217 readily identified species (these exclude the 23 most common readily identified species) they collected via random sampling at the Rothamsted Experimental Station in England. Figure 10 plots the numbers of individual butterflies collected from the 217 species when sorted in rank order (so that the numbers are nonincreasing).

To build a model appropriate for Figure 10, we must include a permutation of the bins as a parameter, since we have sorted the data (see Subsection 4.2 for further discussion of sorting and permutations). We take the model to be

$$p_k(\theta_0, \theta_1) = A_{\theta_1} \frac{(\theta_1)^{\theta_0(k)}}{\sqrt{\theta_0(k) + 23}} \quad (19)$$

for $k = 1, 2, \dots, 216, 217$, where θ_0 is a permutation of the integers $1, 2, \dots, 216, 217$, the parameter θ_1 is a positive real number less than 1, and

$$A_{\theta_1} = \frac{1}{\sum_{k=1}^{217} (\theta_1)^k / \sqrt{k + 23}}; \quad (20)$$

we estimate θ_0 and θ_1 via maximum-likelihood methods (thus obtaining θ_0 by sorting the frequencies into nonincreasing order). Please note that this model is not very carefully chosen — the model is just a truncated geometric distribution weighted by the nonsingular function $1/\sqrt{\theta_0(k) + 23}$, with 23 being the number of common species omitted from the collection. More complicated models may fit better.

The significance levels calculated via 4,000,000 Monte-Carlo simulations are

- χ^2 : .0050
- G^2 : .349
- Freeman-Tukey: .951
- root-mean-square: .00002, i.e., 2E-5

As Figure 10 indicates, the discrepancy between the empirical data and the model is substantial, and, given the large number of draws (5300), cannot be due solely to random fluctuations. The log-likelihood-ratio (G^2) and Freeman-Tukey statistics are unable to detect this discrepancy, while the root-mean-square easily determines that the discrepancy is very highly significant.

4.8 A modified geometric law for religious affiliations

The Pew Forum on Religion and Public Life (a project of the Pew Research Center) recently released [12] — a report on the religious affiliations of Americans — based on a 2007 survey of 35,556 individuals from the continental United States (the full report includes data on Alaska and Hawaii, too, but we chose not to incorporate these). We analyze the identifications reported in the variable “DENOM” from the publicly available data set (“DENOM”




Figure 10: Numbers of specimens (the dots) from 217 species of butterflies (one bin per species), and the best-fit distribution (the lines)

provides the most detailed information on religious affiliations). The 35,556 randomly selected Americans reported affiliations with 372 different religious denominations (of course, it is unlikely that the sample included members from every denomination to which Americans belong; there are undoubtedly more than 372 denominations). Figure 11 plots the numbers of surveyed individuals associated with the various religious denominations when sorted in rank order (so that the numbers are nonincreasing).

To build a model appropriate for Figure 11, we must include a permutation of the bins as a parameter, since we have sorted the data (see Subsection 4.2 for further discussion of sorting and permutations). Furthermore, the tail of the distribution plotted in Figure 11 seems to be more easily modeled than the full distribution. In order to focus the goodness-of-fit test on the tail alone, we can introduce one parameter per bin outside the tail, with the parameter being the probability of drawing the bin under the model. With such a parameter, the model will fit the empirical data exactly in the associated bin — there will be no discrepancy between the data and the model in that bin — so that the bin will not contribute to any goodness-of-fit statistic, aside from altering the number of draws in the remaining bins. To summarize, we need the following parameters: a permutation θ_0 associated with sorting the data, real numbers $\theta_1, \theta_2, \dots, \theta_{54}, \theta_{55}$ specifying the probabilities associated with the first 55 bins in the sorted distribution, and a parameter θ_{56} associated with the model distribution for the tail (which we choose to be a geometric distribution).

Thus, we arrive at the model

$$p_k(\theta_0, \theta_1, \dots, \theta_{55}, \theta_{56}) = \begin{cases} \theta_{\theta_0(k)}, & \theta_0(k) = 1, 2, \dots, 54, 55 \\ (\theta_{56})^{\theta_0(k)-56} (1 - \theta_{56}) (1 - \sum_{j=1}^{55} \theta_j), & \theta_0(k) = 56, 57, 58, \dots \end{cases} \quad (21)$$

for $k = 1, 2, 3, \dots$, where θ_0 is a permutation of the positive integers, and $\theta_1, \theta_2, \dots, \theta_{55}, \theta_{56}$ are real numbers between 0 and 1. While this model may seem complicated at first glance, the estimation of its parameters is actually very simple: first we sort the frequencies into nonincreasing order (thus obtaining θ_0), then we set $\theta_1, \theta_2, \dots, \theta_{54}, \theta_{55}$ to be the 55 greatest numbers of draws divided by the total number (35,556) of draws, and finally we choose θ_{56} to be the base of the geometric distribution which best fits the remaining numbers of draws in the maximum-likelihood sense. The permutation θ_0 lets us sort the data so that the frequencies are in nonincreasing order. The parameters $\theta_1, \theta_2, \dots, \theta_{54}, \theta_{55}$ effectively allow us to ignore the bins with the 55 greatest numbers of draws, as our model fits those bins exactly, by construction. The parameter θ_{56} is the base in the geometric distribution which best fits the tail of the distribution of the data. Figure 11 plots the numbers of surveyed individuals associated with the various religious denominations when sorted in rank order (so that the numbers are nonincreasing), as well as the best-fit model distribution defined in (21). Of the 35,556 total surveyed individuals, 4,050 are not associated with the 55 most popular denominations (that is, 4,050 are not associated with the bins containing the 55 greatest numbers of surveyed individuals).

Since the model defined in (21) involves infinitely many bins, this provides a good opportunity to consider an example of rebinning. Instead of using (21) directly, we rebin so that there are only $n = 340$ bins in all, aggregating the numbers of draws from bins 340, 341, 342, \dots in the original distribution to be the number of draws for bin 340 in the rebinned distribution. We employ the rebinning only for the calculation of the goodness-of-fit statistics; we estimate all parameters, $\theta_0, \theta_1, \dots, \theta_{55}, \theta_{56}$, directly from the data without rebinning, and we generate draws from the estimated model distribution without rebinning when computing the significance levels via Monte-Carlo simulations. (Strictly speaking, for the parameter estimation and Monte-Carlo simulations, we rebin the infinitely many bins down to only 34,000, but for these purposes 34,000 is effectively infinite.)

The significance levels calculated via 1,000,000 Monte-Carlo simulations are then

- χ^2 : .460
- G^2 : .984
- Freeman-Tukey: .992
- root-mean-square: .0011

As Figure 11 indicates, the discrepancy between the empirical data and the model is substantial, and, given the large number of draws, cannot be due solely to random fluctuations. The classical statistics are unable to detect this discrepancy, while the root-mean-square easily determines that the discrepancy is highly significant.





Figure 11: Numbers of surveyed Americans (the dots) identifying with 400 different religious denominations (the bins), and the best-fit distribution (the lines); the fit is perfect for bins 1–55 by definition

5 The power and efficiency of the root-mean-square

In this section, we consider many numerical experiments and models, plotting the numbers of draws required for goodness-of-fit statistics to detect divergence from the models. We consider both fully specified models and parameterized models. To quantify a statistic’s success at detecting discrepancies from the models, we use the formulation of the following remark.

Remark 5.1. We say that a statistic based on given i.i.d. draws “distinguishes” the actual underlying distribution of the draws from the model distribution to mean that the computed confidence level is at least 99% for 99% of 40,000 simulations, with each simulation generating m i.i.d. draws according to the actual distribution. (Recall that a significance level of α is the same as a confidence level of $1 - \alpha$.) We computed the confidence levels by conducting another 40,000 simulations, with each simulation generating m i.i.d. draws according to the model distribution. In Appendix A, we use a weaker notion of “distinguish” — we say that a statistic based on given i.i.d. draws “distinguishes” the actual underlying distribution of the draws from the model distribution to mean that the computed confidence level is at least 95% for 95% of 40,000 simulations, while running simulations and computing confidence levels exactly as for the plots in the present section.

Remark 5.2. To compute the confidence levels for each example in Subsection 5.2, we should in principle calculate the maximum-likelihood estimate $\hat{\theta}$ for each of 40,000 simulations and (for each goodness-of-fit statistic) use these estimates to perform $(40,000)^2$ times the three-step procedure described in Remark 3.3. The computational costs for generating the plots in Subsection 5.2 would then be excessive. Instead, when computing the confidence levels as a function of the value of the statistic under consideration, we calculated $\hat{\theta}$ only once, using as the empirical data 1,000,000 draws from the underlying distribution, and (for each goodness-of-fit statistic) performed 40,000 times the three-step procedure described in Remark 3.3, using the single value of $\hat{\theta}$. The parameter estimates did not vary much over the 40,000 simulations, so approximating the confidence levels thus is accurate. Furthermore, when the parameter is just a permutation, as in Subsections 5.2.8 and 5.2.9, the “approximation” described in the present remark is exactly equivalent to recomputing the confidence levels 40,000 times — we are not making any approximation at all. Please note that we did recalculate the maximum-likelihood estimate $\hat{\theta}$ (and $\tilde{\theta}$ from Remark 3.3) for each of 40,000 simulations when computing the values of the statistics for the simulation; however, when calculating the confidence levels as a function of the values of the statistics, we always drew from the model distribution associated with the same value of the parameter.

Remark 5.3. The root-mean-square statistic is not very sensitive to relative discrepancies between the model and actual distributions in bins whose associated model probabilities are small. When sensitivity in these bins is desirable, we recommend using both the root-mean-square statistic and an asymptotically equivalent variation of χ^2 , such as the log-likelihood-ratio or “ G^2 ” test; see, for example, [17].

5.1 Examples without parameter estimation

5.1.1 A simple, illustrative example

Let us first specify the model distribution to be

$$p_1 = \frac{1}{4}, \quad (22)$$

$$p_2 = \frac{1}{4}, \quad (23)$$

and

$$p_k = \frac{1}{2n-4} \quad (24)$$

for $k = 3, 4, \dots, n-1, n$. We consider m i.i.d. draws from the distribution

$$\tilde{p}_1 = \frac{3}{8}, \quad (25)$$

$$\tilde{p}_2 = \frac{1}{8}, \quad (26)$$

and

$$\tilde{p}_k = p_k \quad (27)$$

for $k = 3, 4, \dots, n-1, n$, where $p_3, p_4, \dots, p_{n-1}, p_n$ are the same as in (24).

Figure 12 plots the percentage of 40,000 simulations, each generating 200 i.i.d. draws according to the actual distribution defined in (25)–(27), that are successfully detected as not arising from the model distribution at the 1% significance level (meaning that the associated statistic for the simulation yields a confidence level of 99% or greater). We computed the significance levels by conducting 40,000 simulations, each generating 200 i.i.d. draws according to the model distribution defined in (22)–(24). Figure 12 shows that the root-mean-square is successful in at least 99% of the simulations, while the classical χ^2 statistic fails often, succeeding in less than 80% of the simulations for $n = 16$, and less than 5% for $n \geq 256$.

Figure 13 plots the number m of draws required to distinguish the actual distribution defined in (25)–(27) from the model distribution defined in (22)–(24). Remark 5.1 above specifies what we mean by “distinguish.” Figure 13 shows that the root-mean-square requires only about $m = 185$ draws for any number n of bins, while the classical χ^2 statistic requires 90% more draws for $n = 16$, and greater than 300% more for $n \geq 128$. Furthermore, the classical χ^2 statistic requires increasingly many draws as the number n of bins increases, unlike the root-mean-square.




Figure 12: First example, with $m = 200$ draws; see Subsection 5.1.1.




Figure 13: First example (statistical “efficiency”); see Subsection 5.1.1.




Figure 14: Second example; see Subsection 5.1.2.

5.1.2 Truncated power-laws

Next, let us specify the model distribution to be

$$p_k = \frac{C_1}{k} \quad (28)$$

for $k = 1, 2, \dots, n-1, n$, where

$$C_1 = \frac{1}{\sum_{k=1}^n 1/k}. \quad (29)$$

We consider m i.i.d. draws from the distribution

$$\tilde{p}_k = \frac{C_2}{k^2} \quad (30)$$

for $k = 1, 2, \dots, n-1, n$, where

$$C_2 = \frac{1}{\sum_{k=1}^n 1/k^2}. \quad (31)$$

Figure 14 plots the number m of draws required to distinguish the actual distribution defined in (30) and (31) from the model distribution defined in (28) and (29). Remark 5.1 above specifies what we mean by “distinguish.” Figure 14 shows that the classical χ^2 statistic requires increasingly many draws as the number n of bins increases, while the root-mean-square exhibits the opposite behavior.




Figure 15: Third example; see Subsection 5.1.3.

5.1.3 Additional truncated power-laws

Let us again specify the model distribution to be

$$p_k = \frac{C_1}{k} \quad (32)$$

for $k = 1, 2, \dots, n-1, n$, where

$$C_1 = \frac{1}{\sum_{k=1}^n 1/k}. \quad (33)$$

We now consider m i.i.d. draws from the distribution

$$\tilde{p}_k = \frac{C_{1/2}}{\sqrt{k}} \quad (34)$$

for $k = 1, 2, \dots, n-1, n$, where

$$C_{1/2} = \frac{1}{\sum_{k=1}^n 1/\sqrt{k}}. \quad (35)$$

Figure 15 plots the number m of draws required to distinguish the actual distribution defined in (34) and (35) from the model distribution defined in (32) and (33). Remark 5.1 above specifies what we mean by “distinguish.” The root-mean-square is not uniformly more powerful than the other statistics in this example; see Remark 5.3 at the beginning of the present section.




Figure 16: Fourth example; see Subsection 5.1.4.

5.1.4 Additional truncated power-laws, reversed

Let us next specify the model distribution to be

$$p_k = \frac{C_{1/2}}{\sqrt{k}} \quad (36)$$

for $k = 1, 2, \dots, n-1, n$, where

$$C_{1/2} = \frac{1}{\sum_{k=1}^n 1/\sqrt{k}}. \quad (37)$$

We now consider m i.i.d. draws from the distribution

$$\tilde{p}_k = \frac{C_1}{k} \quad (38)$$

for $k = 1, 2, \dots, n-1, n$, where

$$C_1 = \frac{1}{\sum_{k=1}^n 1/k}. \quad (39)$$

Figure 16 plots the number m of draws required to distinguish the actual distribution defined in (38) and (39) from the model distribution defined in (36) and (37). Remark 5.1 above specifies what we mean by “distinguish.” Figure 16 shows that the classical χ^2 statistic requires many times more draws than the root-mean-square, as the number n of bins increases.




Figure 17: Fifth example; see Subsection 5.1.5.

5.1.5 A final example with fully specified truncated power-laws

Let us next specify the model distribution to be

$$p_k = \frac{C_2}{k^2} \quad (40)$$

for $k = 1, 2, \dots, n-1, n$, where

$$C_2 = \frac{1}{\sum_{k=1}^n 1/k^2}. \quad (41)$$

We again consider m i.i.d. draws from the distribution

$$\tilde{p}_k = \frac{C_1}{k} \quad (42)$$

for $k = 1, 2, \dots, n-1, n$, where

$$C_1 = \frac{1}{\sum_{k=1}^n 1/k}. \quad (43)$$

Figure 17 plots the number m of draws required to distinguish the actual distribution defined in (42) and (43) from the model distribution defined in (40) and (41). Remark 5.1 above specifies what we mean by “distinguish.” The root-mean-square is not uniformly more powerful than the other statistics in this example; see Remark 5.3 at the beginning of the present section.




Figure 18: Sixth example; see Subsection 5.1.6.

5.1.6 Modified Poisson distributions

Let us specify the model distribution to be the (truncated) Poisson distribution

$$p_k = \frac{B_{3n/8} \left(\frac{3n}{8}\right)^{k-1}}{(k-1)!} \quad (44)$$

for $k = 1, 2, \dots, n-1, n$, where

$$B_{3n/8} = \frac{1}{\sum_{k=1}^n \left(\frac{3n}{8}\right)^{k-1} / (k-1)!}. \quad (45)$$

We consider m i.i.d. draws from the distribution

$$\tilde{p}_{(3n/8)-1} = S/10, \quad (46)$$

$$\tilde{p}_{3n/8} = 4S/5, \quad (47)$$

$$\tilde{p}_{(3n/8)+1} = S/10, \quad (48)$$

$$S = p_{(3n/8)-1} + p_{3n/8} + p_{(3n/8)+1}, \quad (49)$$

$$\tilde{p}_k = p_k \quad (50)$$

for the remaining values of k (for $k = 1, 2, \dots, \frac{3n}{8} - 3, \frac{3n}{8} - 2$ and $k = \frac{3n}{8} + 2, \frac{3n}{8} + 3, \dots, n-1, n$), where p_k is defined in (44).

Figure 18 plots the number m of draws required to distinguish the actual distribution defined in (46)–(50) from the model distribution defined in (44) and (45). Remark 5.1 above specifies what we mean by “distinguish.”




Figure 19: Seventh example; see Subsection 5.1.7.

5.1.7 A truncated power-law and a truncated geometric distribution

Let us finally specify the model distribution to be

$$p_k = \frac{C_1}{k} \quad (51)$$

for $k = 1, 2, \dots, 99, 100$, where

$$C_1 = \frac{1}{\sum_{k=1}^{100} 1/k}. \quad (52)$$

We consider m i.i.d. draws from the (truncated) geometric distribution

$$\tilde{p}_k = c_t t^k \quad (53)$$

for $k = 1, 2, \dots, 99, 100$, where

$$c_t = \frac{1}{\sum_{k=1}^{100} t^k}; \quad (54)$$

Figure 19 considers several values for t .

Figure 19 plots the number m of draws required to distinguish the actual distribution defined in (53) and (54) from the model distribution defined in (51) and (52). Remark 5.1 above specifies what we mean by “distinguish.” See the next section, Subsection 5.2.1, for a similar example, this time involving parameter estimation.




Figure 20: First example; see Subsection 5.2.1.

5.2 Examples with parameter estimation

5.2.1 A truncated power-law and a truncated geometric distribution

We turn now to models involving parameter estimation (for details, see [16]). Let us specify the model distribution to be the Zipf distribution

$$p_k(\theta) = \frac{C_\theta}{k^\theta} \quad (55)$$

for $k = 1, 2, \dots, 99, 100$, where

$$C_\theta = \frac{1}{\sum_{k=1}^{100} 1/k^\theta}; \quad (56)$$

we estimate the parameter θ via maximum-likelihood methods. We consider m i.i.d. draws from the (truncated) geometric distribution

$$\tilde{p}_k = c_t t^k \quad (57)$$

for $k = 1, 2, \dots, 99, 100$, where

$$c_t = \frac{1}{\sum_{k=1}^{100} t^k}; \quad (58)$$

Figure 20 considers several values for t .

Figure 20 plots the number m of draws required to distinguish the actual distribution defined in (57) and (58) from the model distribution defined in (55) and (56), estimating the parameter θ in (55) and (56) via maximum-likelihood methods. Remark 5.1 above specifies what we mean by “distinguish.”




Figure 21: Second example; see Subsection 5.2.2.

5.2.2 A rebinned geometric distribution and a truncated power-law

Let us specify the model distribution to be

$$p_k(\theta) = \theta^{k-1}(1 - \theta) \quad (59)$$

for $k = 1, 2, \dots, 98, 99$, and

$$p_{100}(\theta) = \theta^{99}; \quad (60)$$

we estimate the parameter θ via maximum-likelihood methods. We consider m i.i.d. draws from the Zipf distribution

$$\tilde{p}_k = \frac{C_t}{k^t} \quad (61)$$

for $k = 1, 2, \dots, 99, 100$, where

$$C_t = \frac{1}{\sum_{k=1}^{100} 1/k^t}; \quad (62)$$

Figure 21 considers several values for t .

Figure 21 plots the number m of draws required to distinguish the actual distribution defined in (61) and (62) from the model distribution defined in (59) and (60), estimating the parameter θ in (59) and (60) via maximum-likelihood methods. Remark 5.1 above specifies what we mean by “distinguish.”




Figure 22: Third example; see Subsection 5.2.3.

5.2.3 Truncated shifted Poisson distributions

Let us specify the model distribution to be the (truncated) Poisson distribution

$$p_k(\theta) = \frac{B_\theta \theta^{k-1}}{(k-1)!} \quad (63)$$

for $k = 1, 2, \dots, 20, 21$, where

$$B_\theta = \frac{1}{\sum_{k=1}^{21} \theta^{k-1}/(k-1)!}; \quad (64)$$

we estimate the parameter θ via maximum-likelihood methods. We consider m i.i.d. draws from the distribution

$$\tilde{p}_k = \frac{\tilde{B}_t 5^{k-1+t}}{(k-1+t)!} \quad (65)$$

for $k = 1, 2, \dots, 20, 21$, where

$$\tilde{B}_t = \frac{1}{\sum_{k=1}^{21} 5^{k-1+t}/(k-1+t)!}; \quad (66)$$

Figure 22 considers several values for t . Clearly, $\tilde{p}_k = p_k(5)$ for $k = 1, 2, \dots, 20, 21$, if $t = 0$.

Figure 22 plots the number m of draws required to distinguish the actual distribution defined in (65) and (66) from the model distribution defined in (63) and (64), estimating the parameter θ in (63) and (64) via maximum-likelihood methods. Remark 5.1 above specifies what we mean by “distinguish.”




Figure 23: Fourth example; see Subsection 5.2.4.

5.2.4 Modified Poisson distributions

Let us specify the model distribution to be the (truncated) Poisson distribution

$$p_k(\theta) = \frac{B_\theta \theta^{k-1}}{(k-1)!} \quad (67)$$

for $k = 1, 2, \dots, n-1, n$, where

$$B_\theta = \frac{1}{\sum_{k=1}^n \theta^{k-1}/(k-1)!}; \quad (68)$$

we estimate the parameter θ via maximum-likelihood methods. We consider m i.i.d. draws from the distribution

$$\tilde{p}_{(3n/8)-1} = S/10, \quad (69)$$

$$\tilde{p}_{3n/8} = 4S/5, \quad (70)$$

$$\tilde{p}_{(3n/8)+1} = S/10, \quad (71)$$

$$S = p_{(3n/8)-1}(3n/8) + p_{3n/8}(3n/8) + p_{(3n/8)+1}(3n/8), \quad (72)$$

and

$$\tilde{p}_k = p_k(3n/8) \quad (73)$$

for the remaining values of k (for $k = 1, 2, \dots, \frac{3n}{8} - 3, \frac{3n}{8} - 2$ and $k = \frac{3n}{8} + 2, \frac{3n}{8} + 3, \dots, n-1, n$), where p_k is defined in (67).

Figure 23 plots the number m of draws required to distinguish the actual distribution defined in (69)–(73) from the model distribution defined in (67) and (68), estimating the parameter θ in (67) and (68) via maximum-likelihood methods. Remark 5.1 above specifies what we mean by “distinguish.”




Figure 24: Fifth example; see Subsection 5.2.5.

5.2.5 An example with a uniform tail

Let us specify the model distribution to be

$$p_1(\theta) = \theta, \quad (74)$$

$$p_2(\theta) = \frac{1}{2} - \theta, \quad (75)$$

and

$$p_k(\theta) = \frac{1}{2n-4} \quad (76)$$

for $k = 3, 4, \dots, n-1, n$; we estimate the parameter θ via maximum-likelihood methods. We consider m i.i.d. draws from the distribution

$$\tilde{p}_1 = \frac{3}{8}, \quad (77)$$

$$\tilde{p}_2 = \frac{3}{8}, \quad (78)$$

and

$$\tilde{p}_k = \frac{1}{4n-8} \quad (79)$$

for $k = 3, 4, \dots, n-1, n$.

Figure 24 plots the number m of draws required to distinguish the actual distribution defined in (77)–(79) from the model distribution defined in (74)–(76), estimating the parameter θ in (74)–(76) via maximum-likelihood methods. Remark 5.1 above specifies what we mean by “distinguish.”




Figure 25: Sixth example; see Subsection 5.2.6.

5.2.6 Another example with a uniform tail

Let us specify the model distribution to be

$$p_1(\theta) = \theta, \quad (80)$$

$$p_2(\theta) = \theta, \quad (81)$$

$$p_3(\theta) = \frac{1}{2} - 2\theta, \quad (82)$$

$$p_k(\theta) = \frac{1}{2n-6} \quad (83)$$

for $k = 4, 5, \dots, n-1, n$; we estimate the parameter θ via maximum-likelihood methods. We consider m i.i.d. draws from the distribution

$$\tilde{p}_1 = \frac{1}{4}, \quad (84)$$

$$\tilde{p}_2 = \frac{1}{8}, \quad (85)$$

$$\tilde{p}_3 = \frac{1}{8}, \quad (86)$$

$$\tilde{p}_k = \frac{1}{2n-6} \quad (87)$$

for $k = 4, 5, \dots, n-1, n$.

Figure 25 plots the number m of draws required to distinguish the actual distribution defined in (84)–(87) from the model distribution defined in (80)–(83), estimating the parameter θ in (80)–(83) via maximum-likelihood methods. Remark 5.1 above specifies what we mean by “distinguish.”




Figure 26: Seventh example; see Subsection 5.2.7.

5.2.7 A model with an integer-valued parameter

Let us specify the model distribution to be

$$p_k(\theta) = \frac{1}{2\theta} \quad (88)$$

for $k = 1, 2, \dots, \theta - 1, \theta$, and

$$p_k(\theta) = \frac{1}{2(n - \theta)} \quad (89)$$

for $k = \theta + 1, \theta + 2, \dots, n - 1, n$; we estimate the parameter θ via maximum-likelihood methods. We consider m i.i.d. draws from the distribution

$$\tilde{p}_1 = \frac{1}{4}, \quad (90)$$

$$\tilde{p}_2 = \frac{1}{4}, \quad (91)$$

$$\tilde{p}_3 = \frac{1}{4}, \quad (92)$$

and

$$\tilde{p}_k = \frac{1}{4n - 12} \quad (93)$$

for $k = 4, 5, \dots, n - 1, n$.

Figure 26 plots the number m of draws required to distinguish the actual distribution defined in (90)–(93) from the model distribution defined in (88) and (89), estimating the parameter θ in (88) and (89) via maximum-likelihood methods. Remark 5.1 above specifies what we mean by “distinguish.”




Figure 27: Eighth example; see Subsection 5.2.8.

5.2.8 Truncated power-laws parameterized with a permutation

Let us specify the model to be the Zipf distribution

$$p_k(\theta) = \frac{C_1}{\theta(k)} \quad (94)$$

for $k = 1, 2, \dots, n-1, n$, where θ is a permutation of the integers $1, 2, \dots, n-1, n$, and

$$C_1 = \frac{1}{\sum_{k=1}^n 1/k}; \quad (95)$$

we estimate the permutation θ via maximum-likelihood methods, that is, by sorting the frequencies: first we choose k_1 to be the number of a bin containing the greatest number of draws among all n bins, then we choose k_2 to be the number of a bin containing the greatest number of draws among the remaining $n-1$ bins, then we choose k_3 to be the number of a bin containing the greatest among the remaining $n-2$ bins, and so on, and finally we find θ such that $\theta(k_1) = 1, \theta(k_2) = 2, \dots, \theta(k_{n-1}) = n-1, \theta(k_n) = n$.

We consider m i.i.d. draws from the distribution

$$\tilde{p}_k = \frac{C_2}{k^2} \quad (96)$$

for $k = 1, 2, \dots, n-1, n$, where

$$C_2 = \frac{1}{\sum_{k=1}^n 1/k^2}. \quad (97)$$

Figure 27 plots the number m of draws required to distinguish the actual distribution defined in (96) and (97) from the model distribution defined in (94) and (95), estimating the parameter θ in (94) via maximum-likelihood methods (that is, by sorting). Remark 5.1 above specifies what we mean by “distinguish.”




Figure 28: Ninth example; see Subsection 5.2.9.

5.2.9 Another model parameterized with a permutation

Let us specify the model distribution to be

$$p_k(\theta) = \begin{cases} 3/8, & \theta(k) = 1 \\ 1/8, & \theta(k) = 2 \\ 1/(2n-4), & \theta(k) = 3, 4, \dots, n-1, \text{ or } n \end{cases} \quad (98)$$

for $k = 1, 2, \dots, n-1, n$, where θ is a permutation of the integers $1, 2, \dots, n-1, n$; we estimate the permutation θ via maximum-likelihood methods, that is, by sorting the frequencies: first we choose k_1 to be the number of a bin containing the greatest number of draws among all n bins, then we choose k_2 to be the number of a bin containing the greatest number of draws among the remaining $n-1$ bins, then we choose k_3 to be the number of a bin containing the greatest among the remaining $n-2$ bins, and so on, and finally we find θ such that $\theta(k_1) = 1, \theta(k_2) = 2, \dots, \theta(k_{n-1}) = n-1, \theta(k_n) = n$.

We consider m i.i.d. draws from the distribution

$$\tilde{p}_1 = 1/4, \quad (99)$$

$$\tilde{p}_2 = 1/4, \quad (100)$$

and

$$\tilde{p}_k = 1/(2n-4) \quad (101)$$

for $k = 3, 4, \dots, n-1, n$.

Figure 28 plots the number m of draws required to distinguish the actual distribution defined in (99)–(101) from the model distribution defined in (98), estimating the parameter θ in (98) via maximum-likelihood methods (that is, by sorting). Remark 5.1 above specifies what we mean by “distinguish.”




Figure 29: Tenth example; see Subsection 5.2.10.

5.2.10 A model with two parameters

For the final example, let us specify the model distribution to be

$$p_1(\theta_1, \theta_2) = \theta_1, \quad (102)$$

$$p_2(\theta_1, \theta_2) = \theta_1, \quad (103)$$

$$p_3(\theta_1, \theta_2) = \theta_2, \quad (104)$$

$$p_4(\theta_1, \theta_2) = \theta_2, \quad (105)$$

and

$$p_k(\theta_1, \theta_2) = \frac{1 - 2\theta_1 - 2\theta_2}{n - 4} \quad (106)$$

for $k = 5, 6, \dots, n - 1, n$; we estimate the parameters θ_1 and θ_2 via maximum-likelihood methods. We consider m i.i.d. draws from the distribution

$$\tilde{p}_1 = \frac{9}{32}, \quad (107)$$

$$\tilde{p}_2 = \frac{3}{32}, \quad (108)$$

$$\tilde{p}_3 = \frac{3}{32}, \quad (109)$$

$$\tilde{p}_4 = \frac{1}{32}, \quad (110)$$

and

$$\tilde{p}_k = \frac{1}{2n-8} \quad (111)$$

for $k = 5, 6, \dots, n-1, n$.

Figure 29 plots the number m of draws required to distinguish the actual distribution defined in (107)–(111) from the model distribution defined in (102)–(106), estimating the parameters θ_1 and θ_2 in (102)–(106) via maximum-likelihood methods. Remark 5.1 above specifies what we mean by “distinguish.”

Acknowledgements

We would like to thank Alex Barnett, Gérard Ben Arous, James Berger, Tony Cai, Sourav Chatterjee, Ingrid Daubechies, Jianqing Fan, Andrew Gelman, Leslie Greengard, Peter W. Jones, Michael O’Neil, Ron Peled, Vladimir Rokhlin, Amit Singer, Joel Tropp, Larry Wasserman, and Douglas A. Wolfe. We would also like to thank Jiayang Gao for her many observations, which include pointing out the identity in (4), showing that the Freeman-Tukey/Hellinger-distance statistic is just a weighted version of the root-mean-square.

A Additional plots of power and efficiency

For each plot in Section 5, this appendix provides a corresponding plot based on a confidence level of 95% (that is, a significance level of 5%), rather than a confidence level of 99% (that is, a significance level of 1%). In this appendix Figures 31–47 set the probabilities of false positives and false negatives both to be 5% in order to determine the required number m of draws, whereas in Section 5 above Figures 13–29 set the probabilities of false positives and false negatives both to be 1% (see Remark 5.1). Similarly, a rejection is deemed successful for Figure 30 at the 5% significance level (or better), whereas a rejection is deemed successful for Figure 12 only at the stricter 1% significance level (or better).




Figure 30: First example, with $m = 100$ draws; see Subsection 5.1.1.




Figure 31: First example (statistical “efficiency”); see Subsection 5.1.1.




Figure 32: Second example; see Subsection 5.1.2.




Figure 33: Third example; see Subsection 5.1.3.




Figure 34: Fourth example; see Subsection 5.1.4.




Figure 35: Fifth example; see Subsection 5.1.5.




Figure 36: Sixth example; see Subsection 5.1.6.




Figure 37: Seventh example; see Subsection 5.1.7.




Figure 38: First example; see Subsection 5.2.1.




Figure 39: Second example; see Subsection 5.2.2.




Figure 40: Third example; see Subsection 5.2.3.




Figure 41: Fourth example; see Subsection 5.2.4.




Figure 42: Fifth example; see Subsection 5.2.5.




Figure 43: Sixth example; see Subsection 5.2.6.




Figure 44: Seventh example; see Subsection 5.2.7.




Figure 45: Eighth example; see Subsection 5.2.8.




Figure 46: Ninth example; see Subsection 5.2.9.




Figure 47: Tenth example; see Subsection 5.2.10.

B Convergence to asymptotic levels

In this appendix, we investigate the convergence rates of significance levels to their asymptotic values in the limit of large numbers of draws. We take all draws directly from the model distributions, and focus on models with real-valued parameters. Needless to say, the model parameters are almost surely known exactly in the limit of large numbers of draws. Maximum-likelihood estimates of the parameters converge to the actual values relatively fast in all examples considered below; the significance levels for the root-mean-square converge as fast or faster than those for the classical statistics from the Cressie-Read power-divergence family (the classical statistics are χ^2 , the log-likelihood-ratio G^2 , and the Freeman-Tukey/Hellinger distance).

Figures 48–51 plot the exact significance level versus the level in the limit of large numbers of draws (we computed the asymptotic levels via the method detailed in [16]). The exact significance level is the estimate obtained via $\ell = 800,000$ Monte-Carlo simulations, with each simulation generating m draws according to the model distribution for the values of the parameters specified below. The significance levels obtained via simulations include error bars whose heights (top to bottom) are about twice the standard errors of the estimated levels; we used Remark 3.4 to estimate the standard errors. Please note that, as the number m of draws increases, the plotted traces converge to the straight line through the origin of unit slope (as they should).

Figure 48 plots the exact significance level (estimated via simulation) versus the level in the limit of large numbers of draws, for the model distribution

$$p_k(\theta) = \theta^{k-1}(1 - \theta) \quad (112)$$

for $k = 1, 2, \dots, 8, 9$, and

$$p_{10}(\theta) = \theta^9; \quad (113)$$

we estimate the parameter θ for the goodness-of-fit statistics via maximum-likelihood methods, using $\theta = 7/10$ in the generation of the m i.i.d. draws for each of the 800,000 simulations. Of the four statistics considered, the root-mean-square clearly converges the fastest, as the number m of draws increases.

Figure 49 plots the exact significance level (estimated via simulation) versus the level in the limit of large numbers of draws, taking the model to be the Zipf distribution

$$p_k(\theta) = \frac{C_\theta}{k^\theta} \quad (114)$$

for $k = 1, 2, \dots, 9, 10$, where

$$C_\theta = \frac{1}{\sum_{k=1}^{10} 1/k^\theta}; \quad (115)$$

we estimate the parameter θ for the goodness-of-fit statistics via maximum-likelihood methods, using $\theta = 7/2$ in the generation of the m i.i.d. draws for each of the 800,000 simulations. Of the four statistics considered, the root-mean-square converges by far the fastest.

Figure 50 plots the exact significance level (estimated via simulation) versus the level in the limit of large numbers of draws, taking the model to be the Zipf distribution

$$p_k(\theta) = \frac{C_\theta}{k^\theta} \quad (116)$$

for $k = 1, 2, \dots, 99, 100$, where

$$C_\theta = \frac{1}{\sum_{k=1}^{100} 1/k^\theta}; \quad (117)$$

we estimate the parameter θ for the goodness-of-fit statistics via maximum-likelihood methods, using $\theta = 5/2$ in the generation of the m i.i.d. draws for each of the 800,000 simulations. Of the four statistics considered, the root-mean-square converges by far the fastest, as the number m of draws increases.

Figure 51 plots the exact significance level (estimated via simulation) versus the level in the limit of large numbers of draws, for the two-parameter model distribution

$$p_k(\theta_1, \theta_2) = \begin{cases} \theta_1, & k = 1, 2 \\ \theta_2, & k = 3, 4 \\ (1 - 2\theta_1 - 2\theta_2)/16, & k = 5, 6, \dots, 19, 20 \end{cases} \quad (118)$$


for $k = 1, 2, \dots, 19, 20$; we estimate the parameters θ_1 and θ_2 for the goodness-of-fit statistics via maximum-likelihood methods, using $\theta_1 = 9/40$ and $\theta_2 = 3/20$ in the generation of the m i.i.d. draws for each of the 800,000 simulations. The root-mean-square and χ^2 statistics behave similarly, converging faster than the log-likelihood-ratio, G^2 , and Freeman-Tukey statistics, as the number m of draws increases.

Remark B.1. It is possible to accelerate the convergence via higher-order asymptotics. Presumably such acceleration is possible for all four statistics considered in this appendix.


Remark B.2. For any family $p(\theta)$ of discrete probability distributions parameterized by a permutation θ that specifies the order of the bins (meaning that there exists a discrete probability distribution r such that $p_k(\theta) = r_{\theta(k)}$ for all k), and for any number m of draws, the confidence levels defined in Remark 3.3 have the following highly desirable property: Suppose that the actual underlying distribution \tilde{p} of the experimental draws is equal to $p(\theta)$ for some (unknown) θ . Suppose further that γ is the confidence level for rejecting (7), calculated for a particular realization of the experiment (the associated significance level is $\alpha = 1 - \gamma$). Consider repeating the same experiment over and over, and calculating the confidence level for each realization, each time using that realization's particular maximum-likelihood estimate of the parameter in the hypothesis (7). Then, the fraction of the confidence levels that are less than γ is equal to γ in the limit of many repetitions of the experiment. This property is a compelling reason to use $d(Q, p(\Theta))$ rather than $d(Q, p(\hat{\theta}))$ in the left-hand side of (9). Also, the procedure of Remark 3.3 can be viewed as a parametric bootstrap approximation (see, for example, [2] and [6]).

In addition, for any family $p(\theta)$ of discrete probability distributions, the confidence levels defined in Remark 3.3 have the following highly desirable property: Suppose that the actual underlying distribution \tilde{p} of the experimental draws is equal to $p(\theta)$ for some (unknown) θ . Consider repeating the experiment over and over, and calculating the confidence level for each realization, each time using that realization's particular maximum-likelihood estimate of the parameter in the hypothesis (7). Then, the resulting confidence levels converge in distribution to the uniform distribution over $(0, 1)$ in the limit of large numbers of draws.


It may be somewhat fortuitous that the scheme in Remark 3.3 has so many favorable properties — see, for example, [1] and [18].




(a) root-mean-square



(b) χ^2




(c) G^2




(d) Freeman-Tukey


Figure 48: Convergence for a (rebinned) geometric distribution with $n = 10$ bins




(a) root-mean-square



(b) χ^2




(c) G^2




(d) Freeman-Tukey


Figure 49: Convergence for a Zipf distribution with $n = 10$ bins




(a) root-mean-square



(b) χ^2




(c) G^2




(d) Freeman-Tukey


Figure 50: Convergence for a Zipf distribution with $n = 100$ bins




(a) root-mean-square



(b) χ^2



(c) G^2



(d) Freeman-Tukey

Figure 51: Convergence for a two-parameter model with $n = 20$ bins

References

- [1] M. J. BAYARRI AND J. O. BERGER, *P-values for composite null models*, J. Amer. Statist. Assoc., 95 (2000), pp. 1127–1142.
- [2] P. J. BICKEL, Y. RITOV, AND T. M. STOKER, *Tailor-made tests for goodness of fit to semiparametric hypotheses*, Ann. Statist., 34 (2006), pp. 721–741.
- [3] A. CLAUSET, C. R. SHALIZI, AND M. E. J. NEWMAN, *Power-law distributions in empirical data*, SIAM Review, 51 (2009), pp. 661–703.
- [4] W. G. COCHRAN, *The χ^2 test of goodness of fit*, Ann. Math. Statist., 23 (1952), pp. 315–345.
- [5] R. B. D’AGOSTINO AND M. A. STEPHENS, *Goodness-of-Fit Techniques*, Marcel Dekker, New York, 1986.
- [6] B. EFRON AND R. TIBSHIRANI, *An Introduction to the Bootstrap*, Chapman & Hall/CRC Press, Boca Raton, Florida, 1993.
- [7] R. C. ELDRIDGE, *Six Thousand Common English Words*, Nabu Press, Charleston, South Carolina, 2010. Reprint of the original 1911 edition; available online at <http://www.archive.org/details/sixthousandcomm00eldrgoog>.
- [8] E. EROSHEVA, E. C. WALTON, AND D. T. TAKEUCHI, *Self-rated health among foreign- and US-born Asian Americans: A test of comparability*, Med. Care, 45 (2007), pp. 80–87.
- [9] R. A. FISHER, A. S. CORBET, AND C. B. WILLIAMS, *The relation between the number of species and the number of individuals in a random sample of an animal population*, J. Animal Ecology, 12 (1943), pp. 42–58.
- [10] A. GELMAN, *A Bayesian formulation of exploratory data analysis and goodness-of-fit testing*, Internat. Stat. Rev., 71 (2003), pp. 369–382.
- [11] S. W. GUO AND E. A. THOMPSON, *Performing the exact test of Hardy-Weinberg proportion for multiple alleles*, Biometrics, 48 (1992), pp. 361–372.
- [12] L. LUGO, S. STENCEL, J. GREEN, G. SMITH, D. COX, A. POND, T. MILLER, E. PODREBARAC, M. RALSTON, A. KOHUT, P. TAYLOR, AND S. KEETER, *U.S. Religious Landscape Survey: Religious Affiliation — Diverse and Dynamic*, Pew Forum on Religion and Public Life, Pew Research Center, Washington, DC, 2008.
- [13] G. MARSAGLIA, *Random number generators*, J. Modern Appl. Stat. Meth., 2 (2003), pp. 2–13.
- [14] K. PEARSON, *On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling*, Philos. Mag. (Ser. 5), 50 (1900), pp. 157–175.

- [15] W. PERKINS, M. TYGERT, AND R. WARD, *Computing the confidence levels for a root-mean-square test of goodness-of-fit*, Appl. Math. Comput., 217 (2011), pp. 9072–9084.
- [16] —, *Computing the confidence levels for a root-mean-square test of goodness-of-fit, II*, Tech. Rep. 1009.2260, arXiv, 2011.
- [17] C. R. RAO, *Karl Pearson chi-square test: The dawn of statistical inference*, in Goodness-of-Fit Tests and Model Validity, C. Huber-Carol, N. Balakrishnan, M. S. Nikulin, and M. Mesbah, eds., Birkhäuser, Boston, 2002, pp. 9–24.
- [18] J. M. ROBINS, A. VAN DER VAART, AND V. VENTURA, *Asymptotic distribution of P-values in composite null models*, J. Amer. Statist. Assoc., 95 (2000), pp. 1143–1156.
- [19] E. RUTHERFORD, H. GEIGER, AND H. BATEMAN, *The probability variations in the distribution of α -particles*, Philos. Mag. (Ser. 6), 20 (1910), pp. 698–707.
- [20] STUDENT, *On the error of counting with a haemocytometer*, Biometrika, 5 (1907), pp. 351–360.
- [21] S. R. S. VARADHAN, M. LEVANDOWSKY, AND N. RUBIN, *Mathematical Statistics*, Lecture Notes Series, Courant Institute of Mathematical Sciences, NYU, New York, 1974.
- [22] L. WASSERMAN, *All of Statistics*, Springer, 2003.
- [23] G. K. ZIPF, *The Psycho-Biology of Language: An Introduction to Dynamic Philology*, Houghton Mifflin, Boston, Massachusetts, 1935.